arXiv:1612.02101v2 [cs.CV] 23 Dec 2016

Mining Pixels: Weakly Supervised Semantic Segmentation Using Image Labels

Qibin Hou[1*] Puneet Kumar Dokania[2*] Daniela Massiceti[2] Yunchao Wei[3]
Ming-Ming Cheng[1] Philip H. S. Torr[2]

[1]CCCE, Nankai University [2] University of Oxford [3]NUS

http://mmcheng.net/MinPixel/

Abstract

We consider the task of learning a classifier for semantic segmentation using weak supervision, in this case, image labels specifying the objects within the image. Our method uses deep convolutional neural networks (CNNs) and adopts an Expectation-Maximization (EM) based approach maintaining the uncertainty on pixel labels. We focus on the following three crucial aspects of the EM based approach: (i) initialization; (ii) latent posterior estimation (E step) and (iii) the parameter update (M step). We show that saliency *and* attention *maps provide good cues to learn an initialization model and allows us to skip the bad local maximum to which EM methods are otherwise traditionally prone. In order to update the parameters, we propose minimizing the combination of the standard softmax loss and the KL divergence between the true latent posterior and the likelihood given by the CNN. We argue that this combination is more robust to wrong predictions made by the expectation step of the EM method. We support this argument with empirical and visual results. We additionally incorporate an approximate intersection-over-union (IoU) term into the loss function for better parameter estimation. Extensive experiments and discussions show that: (i) our method is very simple and intuitive; (ii) requires only image-level labels; and (iii) consistently outperforms other weakly/semi supervised state-of-the-art methods with a very high margin on the PASCAL VOC 2012 dataset.*

1. Introduction

The *semantic segmentation task*, the assignment of a semantic class label to every pixel in an image, has ridden the success wave of deep Convolutional Neural Networks (CNNs), resulting in recent and rapid progress on established segmentation benchmarks [5, 6, 25, 40]. The performance of CNNs, however, is largely dependent on the availability of a large corpus of annotated training data, which for the *semantic segmentation task* is both cost- and time-intensive to acquire. It is estimated that the pixel-level annotation of an image in PASCAL VOC takes on average 4 minutes per image [4]. This is likely a conservative estimate given that it is based on the COCO dataset [22] in which ground-truths are obtained by annotating polygon corners rather than pixels directly. In response, recent years have seen an increased interest in weakly-supervised semantic segmentation [4, 18, 29, 30, 32, 33, 38]. These works differ from fully-supervised cases in that rather than having pixel-level ground-truth segmentations, the supervision available is to some lesser degree, for example, image-level labels [18, 29, 30, 32], bounding boxes [29], or points and scribbles [4, 21, 36].

In this work, we address the *semantic segmentation task* using image labels, which specify the objects within the image. Our motivation for using image labels only is two-fold: (i) the annotation of an image with 20 object classes in PASCAL VOC is estimated to take 20 seconds, which is a factor of at least 12 faster than a pixel-level annotation and is also scalable, and (ii) images with their image labels or tags can easily be downloaded from the Internet, providing a rich and virtually infinite source of training data. For a fair analysis of our proposed method with other existing approaches, we provide a detailed comparison of existing weakly/semi-supervised semantic segmentation methods and their dependencies in Tables 1 and 2, respectively.

The method we adopt, similarly to the weakly-supervised semantic segmentation of [29], takes the form of an Expectation-Maximization (EM) [10, 27] based approach. Broadly speaking, an EM based approach has three key steps: (i) initialization; (ii) latent posterior estimation (E step); and (ii) parameter update (M step). We focus on all of these aspects. In particular, an EM based approach is highly susceptible to a mediocre local maximum and its convergence highly depends on the initialization [15]. To handle this issue, we provide an informed initialization to

*Joint first author

(a) Original Image (b) Saliency and Attention (c) Mined Pixels

Figure 1: An example of pixel mining. Figure (1a): original image with two object categories 'person' and 'cow'. Figure (1b): overlap of the saliency map (green) and attention map (red for 'cow' and blue for 'person'). Figure (1c): mined pixels using the combination of the *saliency* map and the *attention* maps (see Algorithm 1). The red pixels belong to the 'cow' and blue to the 'person' object categories.

the EM algorithm by training an initial model for the *semantic segmentation task* using a set of pixels mined using a combined class-agnostic *saliency* map [23] and class-specific *attention* maps [39]. An example of such mining is shown in Figure 1. We mine these pixels from an easy dataset containing simple images from the *ImageNet* [11] dataset (by simple we refer to images with a single object and a clean background). We use saliency and attention maps based cues to automatically select these images from a large corpus of ImageNet dataset. This differs from the work of [29] in which the EM algorithm is initialized using a network trained for the *classification task* on the same *ImageNet* dataset. With regard to our second focus on the latent posterior estimation (E step), we obtain this distribution by constraining the CNN likelihood using a class labels based prior. This in turn reduces many false positive predictions by redistributing the probability masses to the labels present in the image (including the background) and ignoring labels not present in the image. Our method is general enough to allow for the incorporation of other priors. Finally, with regards to our third focus on the *parameter update* (M step), we propose to minimize the combination of the standard *softmax* loss, and the Kullback-Leibler (KL) [20] divergence between the latent posterior probability distribution (obtained using the E step) and the likelihood given by the CNN. The standard *softmax* loss traditionally used to train CNNs approximates the posterior probability distribution using a Dirac delta function. This assigns the probability of 1 to the ground truth semantic class and 0 to all other classes, and minimizing the *softmax* loss between this Dirac delta posterior distribution and the CNN's likelihood is essentially equivalent to minimizing the KL divergence between these two distributions. We argue that using the above-mentioned Dirac delta makes sense in the fully-supervised setting where a human annotator has assigned semantic labels to pixels with very high confidence. In the weakly-supervised setting, however, where the latent posterior can go completely wrong, it is more suitable (and robust to wrong predictions in the case of confusing classes) to minimize the KL divergence between the true latent posterior distribution obtained using the E step and the likelihood distribution given by the CNN. For example, for a particular pixel, if the posterior distribution assigns the probability of 0.5 to 'cat' and 0.49 to 'dog', then it makes more sense to consider these true probabilities while optimizing the parameters, as opposed to the standard *softmax* based approach in which a probability of 1 is assigned to 'cat' and 0 to 'dog'. In addition to this, in order to obtain better CNN parameters, we add a probabilistic approximation of the Intersection-over-Union (IoU) [1, 9, 28] to the above loss function. On the PASCAL VOC 2012 semantic segmentation benchmark [12], we show that with these modifications we obtain state-of-the-art results in the weakly-supervised *semantic segmentation task* with a very high margin compared to the existing weakly/semi-supervised methods.

2. Related Work

In last few years CNNs have been frequently used to address the weakly and semi-supervised semantic segmentation task. There have been many works with varying levels of partial supervision including image labels [18, 29, 30, 38], annotated points [4], squiggles [21, 36], and bounding boxes [29]. Papandreou et al [29] employ an EM based approach with supervision from image labels and bounding boxes. Their method iterates between inferring a latent segmentation (E step) and optimizing the parameters of a segmentation network (M step) by treating the inferred latents as the ground-truth segmentation. Along a similar vein, [38] present a method that trains an initial network using saliency maps, following which a more powerful network is trained using the output of the initial network. The MIL frameworks of [32] and [31] use fully convolutional networks to learn pixel-level semantic segmentations from only image labels. The image labels alone, however, provide no explicit information about the position of the objects in an image. To address this, localization cues can be used [32, 33], obtained through indirect methods like bottom-up proposal generation (for example, using MCG [3] which

is trained using pixel-segmented images), or saliency- [38] and attention-based [39] mechanisms. Localization cues can also be obtained directly through point- and squiggle-level annotations [4, 21, 36].

Our method is most similar to the EM based approach of [29]. We use *saliency* and *attention* cues to learn a network for a simplified version of the *semantic segmentation task* which allows us to better initialize the EM algorithm. This is in contrast to [29] where a network trained for a classification task is used as initialization. Also different from [29] where the latent posterior is approximated by a Dirac delta function (which we argue is too harsh of a constraint in a weakly-supervised setting), we instead propose to use the combination of the true posterior distribution and the Dirac delta function to learn the parameters.

3. The Semantic Segmentation Task

Consider an image I consisting of a set of pixels $\{y_1, \cdots, y_n\}$. Each pixel can be thought of as a random variable which takes on a value from a discrete semantic label set $\mathcal{L} = \{l_0, l_1, \cdots, l_c\}$, where c is the number of semantic classes (or labels) and l_0 denotes the background class. Under this setting, a semantic segmentation is defined as the assignment of all pixels to their corresponding semantic labels, denoted as $\mathbf{y}$.

CNNs are extensively used to model the parameterized class-conditional likelihood for the *semantic segmentation task*. Specifically, assuming each random variable (or pixel) to be independent, a CNN models the likelihood function of the form $P(\mathbf{y}|I;\theta) = \prod_{m=1}^{n} p(y_m|I;\theta)$, where $p(y_m = l|I;\theta)$ is the *softmax* probability (or the marginal) of assigning label l to the m-th pixel. These probabilities are obtained by applying the *softmax*[1] function to the CNN outputs $f(y_m|I;\theta)$ such that $p(y_m = l|I;\theta) \propto \exp(f(y_m = l|I;\theta))$. Given a training dataset $\mathcal{D} = \{I_i, \mathbf{y}_i\}_{i=1}^{N}$, where I_i and $\mathbf{y}_i$ represent the i-th image and its corresponding ground-truth semantic segmentation, respectively, the log-likelihood is usually maximized by minimizing the cross-entropy loss function using the back-propagation algorithm to obtain the optimal θ, referred to as training the CNN. At test time, for a given image, the learned θ is used to obtain the *softmax* probabilities for each pixel. These probabilities are either post-processed or used directly to assign semantic classes to each pixel.

4. Weakly-Supervised Semantic Segmentation

As mentioned in Section 3, to find the optimal θ for the *semantic segmentation task*, we need a dataset with ground-truth pixel-level semantic labels, obtaining which is a highly time-consuming and expensive task. For example, for a given image, annotating it with 20 object classes takes almost 20 seconds, while pixel-wise segmentation takes almost 239.7 seconds [4], making it highly non-scalable to higher numbers of images and classes. In order to handle this issue, focus has turned toward weakly-supervised methods. Motivated by the above-mentioned facts and the work by [29], we use an Expectation-Maximization (EM) [10, 27] based approach for weakly-supervised semantic segmentation with only image labels. Specifically, let us denote $Z = \mathcal{L} \setminus l_0$ as the set of classes that can be present in an image. We denote a weak dataset as $\mathcal{D}_w = \{I_i, \mathbf{z}_i\}_{i=1}^{N}$, where I_i and $\mathbf{z}_i \subseteq Z$ are the i-th image and corresponding labels present in the image, respectively. The objective thus is to learn an optimal θ using $\mathcal{D}_w$ for the semantic segmentation.

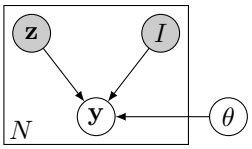

Figure 2: The graphical model. I is the image and $\mathbf{z}$ is the set of labels present in the image. $\mathbf{y}$ is the latent variable (semantic segmentation) and θ is the set of parameters.

4.1. The EM Algorithm

Similarly to [29], we treat the unknown semantic segmentation $\mathbf{y}$ as the latent variable and optimize θ using an EM based approach. Our probabilistic graphical model is of the following form (Figure (2)):

$$P(I, \mathbf{y}, \mathbf{z}; \theta) = P(I)P(\mathbf{y}|I, \mathbf{z}; \theta)P(\mathbf{z}) \tag{1}$$

where we assume that $P(\mathbf{y}|I, \mathbf{z}; \theta) = P(\mathbf{y}|I; \theta)P(\mathbf{y}|\mathbf{z})$. Briefly, to learn θ while maximizing the above joint probability distribution, the three major steps of an EM algorithm are: (i) initialize the parameter θ_t; (ii) E-step: compute the expected complete-data log-likelihood $F(\theta; \theta_t)$, where the expectation is taken over the posterior over the latent variables at a given set of parameters, denoted as $P(\mathbf{y}|I, \mathbf{z}; \theta_t)$ (iii) M-step: update θ by maximizing $F(\theta; \theta_t)$. In what follows, we talk about these steps in detail and show how to learn θ using an EM algorithm.

4.2. Optimizing Parameters

By definition, the expected complete-data log-likelihood is $F(\theta; \theta_t) = \sum_{\mathbf{y}} P(\mathbf{y}|I, \mathbf{z}; \theta_t) \log P(I, \mathbf{y}, \mathbf{z}; \theta)$. In the case of semantic segmentation, the latent space is exponentially large $|\mathcal{L}|^n$, therefore, computing $F(\theta; \theta_t)$ is infeasible. However, as will be shown, the independence assumption over the random variables, namely $P(\mathbf{y}|I; \theta) = \prod_{m=1}^{n} p(y_m|I; \theta)$, allows us to maximize $F(\theta; \theta_t)$ efficiently by decomposition. By using equation (1), the independence assumption over the random variables, the iden-

[1] the *softmax* function is defined as $\sigma(f_k) = \frac{e^{f_k}}{\sum_{j=1}^{c} e^{f_j}}$

tity $\sum_{\mathbf{y}} P(\mathbf{y}|I,\mathbf{z};\theta_t) = 1$, and ignoring the terms independent of θ, $F(\theta;\theta_t)$ can be written in a simplified form as:

$$\bar{F}(\theta;\theta_t) = \sum_{m=1}^{n} \sum_{\mathbf{y}} P(\mathbf{y}|I,\mathbf{z};\theta_t) \log p(y_m|I;\theta) \quad (2)$$

Without loss of generality, we can write $P(\mathbf{y}|I,\mathbf{z};\theta_t) = P(\mathbf{y} \setminus y_m|I,\mathbf{z},y_m;\theta_t)p(y_m|I,\mathbf{z};\theta_t)$, and using the identity $\sum_{\mathbf{y}\setminus y_m} P(\mathbf{y} \setminus y_m|I,\mathbf{z},y_m;\theta_t) = 1$, we obtain:

$$\bar{F}(\theta;\theta_t) = \sum_{m=1}^{n} \sum_{y_m \in \mathcal{L}} p(y_m|I,\mathbf{z};\theta_t) \log p(y_m|I;\theta) \quad (3)$$

Therefore, the M-step parameter update, which is maximizing $\bar{F}(\theta;\theta_t)$ w.r.t. θ, can be written as:

$$\theta_{t+1} = \underset{\theta}{\operatorname{argmax}} \sum_{m=1}^{n} \sum_{y_m \in \mathcal{L}} p(y_m|I,\mathbf{z};\theta_t) \log p(y_m|I;\theta) \quad (4)$$

Based on our assumption mentioned in Section 4.1, the latent posterior distribution can be written as $P(\mathbf{y}|I,\mathbf{z};\theta_t) = P(\mathbf{y}|I;\theta_t)P(\mathbf{y}|\mathbf{z})$, where $P(\mathbf{y}|I;\theta_t)$ is the likelihood obtained using the CNN at a given θ_t. The distribution $P(\mathbf{y}|\mathbf{z})$ can be used to *regularize* the CNN likelihood $P(\mathbf{y}|I;\theta_t)$ by imposing constraints based on the image label information that we have in the dataset. Note that, $P(\mathbf{y}|\mathbf{z})$ is independent of θ and is a task-specific user-defined distribution that depends on the image labels. For example, if we know that there are only two classes in a given training image such as 'cat' and 'person' out of many possible classes, then we would like to push the latent posterior probability $P(\mathbf{y}|I,\mathbf{z};\theta_t)$ of absent classes to zero and increase the probability of the present classes. We can impose this constraint as follows. Let us assume that similar to the likelihood, $P(\mathbf{y}|\mathbf{z})$ also decomposes over the pixels. Therefore, the posterior can be written as $P(\mathbf{y}|I,\mathbf{z};\theta_t) = \prod_{m=1}^{n} p(y_m|I,\mathbf{z};\theta_t)$, where $p(y_m|I,\mathbf{z};\theta_t) \propto p(y_m|I;\theta_t)p(y_m|\mathbf{z})$. Thus, to enforce the above mentioned constraint we use uniform distribution for the classes present in the image including the background (l_0) and a zero probability for the other classes:

$$p(y_m|\mathbf{z}) = \begin{cases} 0, & if\ y_m \notin \mathbf{z} \cup l_0, \\ \frac{1}{|\mathbf{z}|+1}, & otherwise. \end{cases} \quad (5)$$

Practically speaking, in order to avoid normalization issues, we assume $p(y_m|\mathbf{z}) \propto \exp(g(y_m,\mathbf{z}))$ (exponential family distribution), where $g(.)$ is a user-defined function. Thus, the posterior can be written as $p(y_m|I,\mathbf{z};\theta_t) \propto \exp(f(y_m|I;\theta_t) + g(y_m,\mathbf{z}))$. In our case, we use $g(y_m = l,\mathbf{z}) = -\infty$ if the label l is not present in the set $\mathbf{z} \cup l_0$, otherwise, we use $g(y_m = l,\mathbf{z}) = 0$. This is equivalent to obtaining *softmax* probabilities for only those classes (including background) present in the image label set and assign a probability of zero to other classes. The above formulation is generalizable in that $g(.)$ can be used to impose other types of label-dependent constraints.

Cross entropy functions Let us now consider the parameter update problem as defined in Eq. 4, also known as the M step. Solving this problem is equivalent to minimizing the cross entropy or the KL divergence between the latent posterior distribution and the CNN likelihood. Notice that, as opposed to [29], which uses a Dirac delta approximation of the posterior distribution $\hat{p}$, where $\hat{p}(\hat{l}_m) = 1$ at $\hat{l}_m = \operatorname{argmax}_{l \in \mathcal{L}} p(y_m = l|I,\mathbf{z};\theta_t)$ and otherwise zero, we use the true posterior distribution (or the regularized CNN likelihood). We argue that using a Dirac delta distribution is a too-*hard* constraint suitable only for cases when we are very confident about the label assignment by the posterior distribution. For example, in the fully-supervised setting it makes sense to put all the probability mass at the label assigned to the pixel by a human annotator. However, in the weakly-supervised setting where the likelihood can be noisy (mostly seen in the case of confusing classes), it is more suitable to use the true posterior distribution obtained using the combination of the CNN likelihood $p(y_m|I,\theta_t)$ and the class label based prior $p(y_m|\mathbf{z})$. In order to use the best of both worlds, we propose to optimize θ by combining the true posterior distribution and the Dirac delta posterior distribution as follows:

$$\begin{aligned} J_m(I,\mathbf{z},\theta_t;\theta) &= (1-\epsilon) \sum_{y_m \in \mathcal{L}} p(y_m|I,\mathbf{z};\theta_t) \log p(y_m|I;\theta) \\ &+ \epsilon \sum_{y_m \in \mathcal{L}} \hat{p}(y_m) \log p(y_m|I;\theta) \end{aligned} \quad (6)$$

where, given the pixel-wise posterior $p(y_m = l|I,\mathbf{z};\theta)$, several heuristics can be employed to decide the weighting parameter ϵ. We define two such heuristics based on a user-defined hyper-parameter $\eta \in [0,1]$. Our first heuristic, *Max Heuristic*, defines ϵ as follows:

$$\epsilon = \begin{cases} 1, & if\ p_{max} \geq \eta, \\ p_{max}, & otherwise. \end{cases} \quad (7)$$

where, $p_{max} = \max_{l \in \mathcal{L}} p(y_m = l|I,\mathbf{z};\theta)$ is the maximum probability value in the latent posterior distribution. Intuitively, the *Max Heuristic* (Eq. 7) implies that if the posterior is confident (relative to η) in assigning a label to a pixel then use the Dirac delta posterior based cross entropy, otherwise use the weighted combination of the true posterior based cross entropy and the Dirac delta posterior based cross entropy, where the weight is decided by p_{max}. Our second heuristic, *Relative Heuristic*, defines ϵ as follows:

$$\epsilon = \begin{cases} 1, & if\ r \geq \eta, \\ r, & otherwise. \end{cases} \quad (8)$$

Algorithm 1 Mining Pixels

input Image I; label set $\mathbf{z}$; threshold τ

Mask: $M = zeros(n, L+1)$, n and L are the number of pixels and labels, respectively. One extra map is for the background.
$\mathbf{s} \leftarrow SaliencyDetection(I)$ [23]
for each $l \in \mathbf{z}$ **do**
  $\mathbf{a} \leftarrow AttentionNetwork(I, l)$ [39]
  for each pixel m **do**
    if $h(\mathbf{s}(m), \mathbf{a}(m)) \geq \tau$ **then**
      $M(m, l) = 1$
    else if $\mathbf{s}(m) \approx 0$ **then**
      $M(m, L+1) = 1$
    end if
  end for
end for

output M

where, $r = (p^1_{max} - p^2_{max})/p^1_{max}$. Terms p^1_{max} and p^2_{max} are the maximum and second maximum probability values in the latent posterior distribution. Intuitively, if r is higher, it means that the confidence of the most probable label is very high compared to the confidence of the second most probable label. In this case, the Dirac delta based posterior should be preferred, otherwise, the weighted combination should be used where the weight depends on the value of r.

The IoU gain function Along with minimizing the cross entropy losses as shown in the Eq. 6, in order to obtain better parameter estimate, we also maximize the intersection-over-union (IoU) between the posterior distribution and the likelihood using the following probabilistic approximation of the IoU [1, 9, 28]:

$$\mathcal{J}_{IOU}(P(\mathbf{y}|I, \mathbf{z}; \theta_t), P(\mathbf{y}|I; \theta)) \approx \frac{1}{|\mathcal{L}|} \sum_{l \in \mathcal{L}} \frac{\sum_{m=1}^{n} p^t_m(l) p^\theta_m(l)}{\sum_{m=1}^{n} \{p^t_m(l) + p^\theta_m(l) - p^t_m(l) p^\theta_m(l)\}} \tag{9}$$

where, $p^t_m(l) = p(y_m = l|I, \mathbf{z}; \theta_t)$ and $p^\theta_m(l) = p(y_m = l|I; \theta)$. Refer to [9] for further details about Eq. 9.

Overall objective function Combining the cross entropy loss function (Eq. 6) and the IoU gain function (Eq. 9), the M-step parameter update problem can be written as:

$$\theta_{t+1} = \underset{\theta}{\operatorname{argmax}} \sum_{m=1}^{n} J_m(I, \mathbf{z}, \theta_t; \theta) + \mathcal{J}_{IOU} \tag{10}$$

We use a CNN model along with the back-propagation algorithm to optimize the above mentioned objective function.

Algorithm 2 Our Final Algorithm

input Datasets $\mathcal{D}_w$ and $\mathcal{D}_e$; θ_0; η, K; τ

Use $\mathcal{D}_e$, τ, and θ_0 to obtain initialization parameter θ_t using method explained in Section 4.3.
for $k = 1 : K$ **do**
  $\theta \leftarrow \theta_t$
  for each pixel m in $\mathcal{D}_w$ **do**
    Obtain latent posterior: $p_m(y_m|I, \mathbf{z}; \theta) \propto \exp(f_m(y_m|I; \theta) + g(y_m, \mathbf{z}))$
  end for
  Optimize Eq. 10 using CNN to obtain θ_{t+1}. Use η in either Eq. 7 or Eq. 8, depending on the heuristic.
end for

4.3. Skipping Poor Local Maximum: Initial Model

It is well known that if the log-likelihood has several maxima or saddle points, which is true in the case of CNNs, an EM based approach is highly susceptible to mediocre local maxima and its convergence to a local maximum greatly depends on the initialization [15]. Thus, a good initialization is crucial for the success of an EM based algorithm.

We train an *initial* CNN *model* in order to better initialize the EM algorithm. Our approach is based on the following two insights: (i) instead of initializing the EM algorithm with the network parameters trained for the classification task (as done in [29]), we should initialize the algorithm with parameters trained for solving the *semantic segmentation task* but on an easier dataset; and (ii) training a network with a few supervised pixels is capable of guiding the learning process towards a generalizable set of parameters [4].

In order to materialize our first insight we use an easy dataset $\mathcal{D}_e$ containing simple images and their corresponding image labels, selected from the *ImageNet* image classification dataset [11]. The simple images contain mostly centered and clutter-free objects, unlike the challenging PASCAL VOC 2012 [12]. Details of our $\mathcal{D}_e$ are provided in Section 5. In order to train our *initial model*, however, we need pixel-level semantic labels which are not available. To circumvent this problem we follow our second insight and *mine* a set of pixels for each image and use them to train our initial model. We propose to combine a class-agnostic *saliency* map [8, 23] with per-class *attention* maps [39] to *mine* these pixels to which we can assign semantic labels with very high confidence. An example of this approach is shown in Figure (1) and the methodology is summarized in Algorithm 1. The confidence function $h(.,.)$ in Algorithm 1 is a user-defined function which decides whether the pixel should be mined or not.

Intuitively, a *saliency* map gives us the probability of each pixel belonging to any foreground class, and an *attention* map gives us the probability of each pixel belonging to

Dependency	Supervision
Class size	Image labels + Bboxes
Saliency	[34] Image labels
	[17] Bboxes
	[23] Saliency masks
Attention [39]	Image labels
Objectness [2]	Image labels + Bboxes
Localization [41]	Image labels
Superpixels [13]	None
Bbox BING [7]	Bboxes
MCG [3]	Pixel labels
SS [35]	Bboxes
CRF [19]	Pixel labels (parameter cross-val)

Table 2: **Dependency table**. The degrees of supervision required by the dependencies of several weakly-supervised methods for the semantic segmentation task. 'Bboxes' represents bounding boxes.

a given object class. Combining these two maps allows us to retrieve class-specific pixels with high confidence (for details, see Section 5). Once we have the class-specific pixels, we optimize the objective Eq. 10 to obtain the initialization of the EM algorithm. Notice that we only use image labels to mine pixels and assign semantic classes to them. Since the dataset used at this stage is easy, the semantic classes assigned to the mined pixels are less noisy. We, therefore, use a Dirac delta posterior based cross entropy (setting $\eta = 0$ in Eq. 10) for training the *initial model*.

4.4. Overall Weakly-Supervised Algorithm

Combining all the steps, the overall algorithm for our approach is given in Algorithm 2. Datasets $\mathcal{D}_w$ and $\mathcal{D}_e$ are the weak and easy datasets, respectively. The parameter θ_0 is trained for the classification task using *ImageNet* dataset. η is the threshold used in Eq. 7 and Eq. 8. The parameter K is the number of EM iterations and τ is the threshold used in Algorithm 1 to mine pixels.

5. Experimental Results and Comparisons

We show the efficacy of our method on the challenging PASCAL VOC 2012 bechmark for the semantic segmentation task. Our method outperforms all the existing state-of-the-art weakly/semi supervised methods for the semantic segmentation task with a large margin. We present additional results for the purpose of better understanding and analysis in the supplementary material (Section 7).

5.1. Setup

Dataset In order to train our *initial model* (Section 4.3) for the initialization purpose, we select images from the *ImageNet* dataset. We collected $80,000$ images from the *ImageNet* dataset corresponding to the 20 foreground object categories of the PASCAL VOC 2012 segmentation task. In order to retain easy images with less cluttered and almost centered objects, we filter this dataset using simple heuristics based on the image size, *saliency* map and *attention* maps. In the end, we are left with $24,000$ images. This dataset is referred to as the easy dataset $\mathcal{D}_e$. We will make this dataset publicly available. In order to train the EM algorithm we use the PASCAL VOC 2012 dataset that contains $1,464$ training images (train) along with the extra annotations provided by [14], giving a total of $10,582$ images, and combine it with the D_e dataset. We call this the weak dataset D_w.

CNN **Architecture and Parameter Settings** Similar to [18, 29, 38], our implementation is based on the well-known largeFOV DeepLab architecture [6]. We use simple bilinear interpolation to map the downsampled feature maps to the original image size as suggested in [25]. We use the publicly available Caffe toolbox [16] based DeepLab for our implementation purposes. The CNN hyper-parameters used are: weight decay (0.0005), momentum (0.9), and iteration size (10) for gradient accumulation. The learning rate is 0.001 at the beginning and is divided by 10 after every 10 epochs. We use a batch size of 1 and randomly crop the input image to 321×321. Images with width or height less than 321 are padded with the mean pixel values and the corresponding places in the ground-truth are padded with ignore labels to nullify the effect of padding. We flip the input images horizontally, resulting in an augmented set twice of the original one. We train our network for 30K iterations.

Training To train the *initial model* we use the DeepLab architecture [6] initialized with the *ImageNet* trained classification network, or θ_0. We use Algorithm 1 to mine pixels from the easy dataset $\mathcal{D}_e$ using $\tau = 0.5$. Since $\mathcal{D}_e$ contains simple images, we use the 'maximum' as our decision function $h(.,.)$ in Algorithm 1. Notice that the set of mined pixels greatly depends on τ, therefore, it is possible to come up with better $h(.,.)$ to mine less noisy pixels. The *initial model* is then trained using the mined pixels by optimizing Eq. 10 with $\eta = 0$ to obtain θ_t. Finally, we use θ_t to initialize our EM based algorithm with $K = 2$ and use D_w as our training dataset. We noticed that after two EM iterations, the gain in performance was not significant. For the *Max* and *Relative* heuristics, we experimented with $\eta = \{0, 0.1, 0.2, 0.3, 0.4\}$ and $\eta = \{0.05, 0.075, 0.1, 0.15, 0.25\}$, respectively.

5.2. Results and Comparisons

In order to allow for a fair analysis, we provide Table 1 and Table 2 comparing current methods, their dependen-

Method	Dataset	Dependencies	Supervision	CRF [19]	mIoU (Val)	mIoU (Test)
EM Adapt [29]	$D(I,P)$	No	Image labels	✗	–	–
				✓	38.2%	39.6%
CCNN [30]	$D(I,P)$	No	Image labels	✗	33.3%	35.6%
				✓	35.3%	–
		Class size		✗	*40.5*%	*43.3*%
				✓	42.4%	45.1%
SEC [18]	$D(I,P)$	Saliency [34] & Localization [41]	Image labels	✗	*44.3*%	–
				✓	50.7%	51.7%
MIL [32]	$D(I)$	Superpixels [13]	Image labels	✗	36.6%	35.8%
		BBox BING [7]			*37.8*%	*37.0*%
		MCG [3]			42.0%	40.6%
WTP [4]	$D(I,P)$	Objectness [2]	Image labels	–	*32.2*%	–
			Image labels + 1 Point/Class		42.7%	–
			Image labels + 1 Squiggle/Class		49.1%	–
STC [38]	$D(I,P,F)$	Saliency [17]	Image labels	✓	49.8%	51.2%
AugFeed [33]	$D(I,P)$	SS [35]	Image labels	✗	*46.98*%	*47.8*%
				✓	52.62%	52.7%
		MCG [3]		✗	50.41%	50.6%
				✓	54.34%	55.5%
Ours Initial Model	$D(I)$	Saliency [23] & Attention [39]	Image labels	✗	53.53%	54.34%
Ours Max Heu	$D(I,P)$				56.86%	57.67%
Ours Rel Heu	$D(I,P)$				**56.91**%	**57.74**%

Table 1: **Comparison table**. All dependencies, datasets, and degrees of supervision used by most of the existing methods for weakly-supervised semantic segmentation. Table 2 complements this table by providing the degrees of supervision used by the dependencies themselves, which is important to draw a fair comparison. Note that, $D(I)$ represents the *ImageNet* dataset and $D(I,P)$ represents the combination of ImageNet and PASCAL VOC 2012 datasets. Dataset $D(I,P,F)$ contains 41K images from *Flickr* along with $D(I,P)$. Note that the cross validation of CRF hyper-parameters and the training of MCG are performed using a fully supervised pixel-level semantic segmentation dataset. To be fair, our method should only be compared with the numbers shown in italic with underline.

cies, and degrees of supervision. To note, most methods using CRF [19] either as a post-processing step or inside their model are using CRF hyper-parameters which have been cross validated over a pixel-level semantic segmentation dataset. It is, therefore, not fair to categorize them as purely weakly supervised. This is likewise the case for MCG [3] which is trained on a pixel-level semantic segmentation dataset. Regarding the dependencies of our method, our *saliency* network [23] is trained using a dataset of salient region masks. These masks and the maps generated are class-agnostic, therefore, once trained, the network can be used for any semantic object category. There is, therefore, no issue with the scalability and no need to train the *saliency* network again for new object categories. Our second dependency, the *attention* network [39] is trained solely using image-level labels. Please note that we do not use CRF [19] at any stage of our training or testing.

It is evident from the Table 1 that our proposed method consistently outperforms all the existing state-of-the-art methods with a very high margin. A fair comparison of our method would be with AugFeed (augmented feedback) [33] with super-pixel based dependencies while not using CRF. Our method obtains almost 10% better mIoU than AugFeed in both PASCAL val and test datasets. Even if we compare our method with the best existing method without worrying about any sort of dependencies or degrees of supervision they use (AugFeed [33] with MCG and CRF, requires pixel-level annotation), our method is almost 2.6% better in val and 2.2% better in the test datasets. Table 3 shows class-wise performance of our method on both val and test datasets.

We show some visual results in Figure 3 in order to better understand our algorithm. Notice that, as we move from the *initial model* (third column) to the model obtained using the first iteration of our EM based algorithm (fourth column), the semantic segmentation accuracy improves drastically. The same effect can clearly be seen from the first iteration (fourth column) to the second iteration (fifth column) of the algorithm. This clearly shows that the initialization of our method using the initial model is allowing our algorithm to obtain a good local maximum, and thus better parameters for the task. Few experiments revealed that after two iterations of our algorithm the gain in the performance was not significant compared to the computational cost.

Data	Method	bkg	plane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mIoU
Val	*Initial*	87.1	74.7	29.0	69.8	55.8	55.6	73.3	65.2	63.4	15.8	61.5	15.9	60.0	56.4	57.5	53.7	32.9	65.6	23.9	64.6	42.2	53.53
	Max	87.9	72.9	29.2	66.9	58.7	57.4	77.2	69.1	70.8	17.3	64.1	24.4	63.4	60.4	62.4	58.7	37.7	68.5	31.6	71.5	44.0	56.86
	Rel	87.8	72.4	28.7	67.9	58.8	55.8	78.0	69.7	70.2	17.8	63.3	23.2	65.7	60.5	63.1	58.7	40.0	68.2	28.9	70.9	45.5	56.91
Test	*Initial*	87.9	69.2	29.2	74.9	41.7	53.4	70.6	69.6	59.9	18.3	66.1	24.9	62.5	63.3	68.8	55.4	33.7	63.8	18.6	64.3	44.9	54.35
	Max	88.2	70.8	29.5	72.3	45.3	57.3	72.8	72.1	69.0	19.8	64.4	34.6	67.4	63.5	72.4	59.1	45.8	69.3	28.1	64.3	45.2	57.67
	Rel	88.2	69.5	29.7	72.2	45.1	57.3	73.2	72.7	69.3	20.5	65.4	33.5	67.8	64.0	72.3	58.9	45.5	69.8	26.8	63.8	46.8	57.74

Table 3: Experimental results (in %) on PASCAL VOC validation and test sets using our proposed method. *Initial* shows the performance of our *initial model* trained solely using an IMAGENET easy dataset for the initialization of our algorithm. *Max* and *Rel* shows our EM based algorithm using *Max* and *Relative* heuristics defined in Section 4.2, respectively. For the *Relative* heuristic, the results are shown using $\eta = 0.05$ for both iterations of our algorithm. For the *Max* heuristic, results are shown using $\eta = 0.1$ and 0 for the first and the second iterations, respectively. Note that, results are shown for the final iteration of our algorithm.

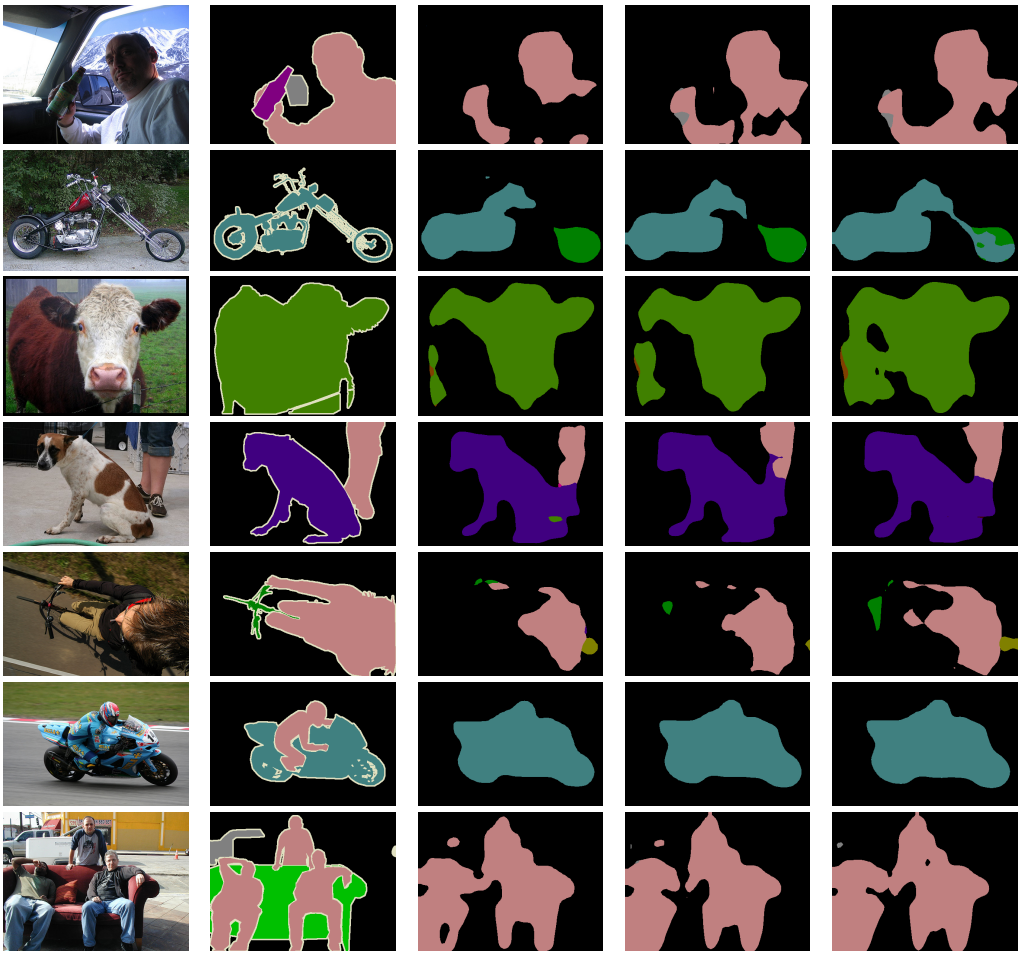

(1) Source (2) Ground truth (3) Initial model (4) EM 1st iteration (5) EM 2nd iteration

Figure 3: Visual results for image labels based semantic segmentation using our proposed EM based method. The two bottom rows are the failure cases. It is clear that as we move from the *Initial model* (3rd column) to the second iteration of the EM algorithm (5th column), the segmentation quality improves drastically. Thus, each EM iteration is learning semantically meaningful information.

6. Conclusions and Future Work

We addressed weakly-supervised semantic segmentation task, where the supervision comes from only image-level labels. We proposed EM based approach and focused on the three key components of the EM algorithm: (i) initialization, (ii) E step and (iii) the M step. Using image labels only and an easy dataset that we automatically generated using the ImageNet images, we learn a set of parameters for the *semantic segmentation task* which provides an informed ini-

tialization for our algorithm. Following this, with each iteration of the EM algorithm, we empirically and qualitatively verify that our method improves the segmentation accuracy on the challenging PASCAL VOC 2012 benchmarks. Furthermore, we showed that our method outperforms the current state-of-the-art methods with very large margin.

In the future, it would be very interesting to further explore the pixel mining aspect of our algorithm as obtaining less noisy pixels is very crucial. In addition to this, it would be very exciting to add one more component to our model that allows us to handle noisy labels. For example, images downloaded from the internet may have more classes than the query key using which they were downloaded.

7. Supplementary

7.1. Analysis

To further explore our method and better understand its successes and failures, here we include additional empirical results on the PASCAL VOC 2012 validation dataset.

Effects of classes Tables 4, 5 and 6 show the class-wise Intersection-over-Union (IOU) for different thresholds η in the *Max* and *Relative Heuristic* settings defined in Section 4.2. For the majority of classes, our method supersedes the results of all other current state-of-the-art methods with equivalent supervision [18, 30, 31, 33]. We observe, however, that our method consistently performs poorly on the 'bike', 'plant' and 'sofa' classes. We attribute this to failures of the attention and saliency networks, which are not being able to detect and attend to these classes well. For example, plants are often not seen as particularly salient objects in an image, and this is similarly the case for bicycles and sofas, since they often occur simultaneously with people (who are more salient-worthy). These failures result in our mining method not producing many good training pixels for these classes, thus affecting the goodness of our initialization parameters for these classes.

Effects of the threshold η on the combined loss As detailed in Section 4.2, the threshold η is a user defined hyper-parameter of the combined loss, controlling the weighted combination of a traditional Dirac delta distribution based cross-entropy loss term (hereafter referred as 'hard') and true posterior distribution cross-entropy loss term (hereafter referred as 'soft') in the overall objective function (see Eq. 6 and Eq. 10). For the *Max Heuristic*, $\eta = 0.1$, for example, enforces that the maximum (latent) probability (p_{max}) must be at least 0.1 (i.e. 10% confident while making a decision) in order to use the 'hard' cross-entropy, otherwise, a weighted combination of both the loss terms should instead be employed. For the *Relative Heuristic*, we instead define a variable r as the relative difference between the highest and second-highest class probability for a pixel. A threshold of $\eta = 0.05$ enforces that the most confident class probability should be at least 5% more confident than the second-most confident class in order to employ the 'hard' cross-entropy, otherwise, the combined loss should be used. Also to note, a threshold of $\eta = 0$ means that only the 'hard' cross-entropy term is employed.

Table 4 shows the class-wise results for different values of η for each heuristic. We considered a range of η which we deemed sensible for the task. For the *Max* and *Relative* heuristics, we experimented with $\eta = \{0, 0.1, 0.2, 0.3, 0.4\}$ and $\eta = \{0.05, 0.075, 0.1, 0.15, 0.25\}$, respectively. Here we are showing results for few settings that gave promising results. The results shown in our main paper (Table 3) are corresponding to the threshold values highlighted using the blue color in Table 4. Table 5 and 6 show the class-wise IOU from our *Initial Model* to our final EM iteration ($K = 2$) corresponding to the best-performing thresholds for the *Max* and *Relative Heuristic*, respectively.

Benefits of the combined loss An important point to note is that when only the 'hard' cross-entropy is used ($\eta = 0$), the results are almost close to when the combined loss is used ($\eta \neq 0$). However, experimentally we observed that the results when using the combined loss are already on-par with the hard loss even after the 1st EM iteration. This suggests that the combined loss encourages the EM algorithm to converge faster. Another interesting point is that the combined loss allow us to define our own heuristics based on the task in hand. Also, the relative weighting can possibly be chosen in a better way to further improve the results.

Heur	Threshold	bkg	plane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mIoU
Max	(0, 0)	87.9	73.0	29.6	66.7	58.8	57.9	77.4	69.0	70.6	16.7	64.7	23.2	64.4	59.3	62.4	58.4	37.5	68.3	31.2	71.0	44.5	56.78
	(0.1, 0)	87.9	72.9	29.2	66.9	58.7	57.4	77.2	69.1	70.8	17.3	64.1	24.4	63.4	60.4	62.4	58.7	37.7	68.5	31.6	71.5	44.0	56.86
	(0.1, 0.1)	87.8	73.4	29.5	67.3	60.0	56.4	77.5	68.4	70.0	16.9	64.1	23.2	64.8	60.7	61.8	57.4	39.5	69.5	29.1	70.6	44.5	56.78
Rel	(0.05, 0)	87.9	72.7	28.5	67.1	59.4	55.5	77.8	69.3	70.6	18.4	62.4	23.9	65.3	58.9	61.3	58.0	39.7	69.8	29.9	71.7	44.7	56.80
	(0.05, 0.05)	87.8	72.4	28.7	67.9	58.8	55.8	78.0	69.7	70.2	17.8	63.3	23.2	65.7	60.5	63.1	58.7	40.0	68.2	28.9	70.9	45.5	56.91

Table 4: Experimental results (in %) on PASCAL VOC validation set using different values of threshold η for the *Max* and *Relative Heuristic* defined in Section 4.2 of the main paper. *(a, b)* in *Threshold* column represents η used in EM *Iteration 1* and EM *Iteration 2*, respectively. The results are shown for the final iteration (second iteration) of our algorithm. Note that the initial model used for the initialization of the EM algorithm was trained using $\eta = 0$.

Data	Stage	bkg	plane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mIoU
val	*Initial*	87.1	74.7	29.0	69.8	55.8	55.6	73.3	65.2	63.4	15.8	61.5	15.9	60.0	56.4	57.5	53.7	32.9	65.6	23.9	64.6	42.2	53.53
	EM-*Iter1*	87.6	74.3	30.5	69.4	58.3	58.7	76.5	66.9	69.6	16.3	64.9	19.4	63.8	61.0	61.3	56.4	36.8	68.4	28.3	69.6	44.1	56.27
	EM-*Iter2*	87.9	72.9	29.2	66.9	58.7	57.4	77.2	69.1	70.8	17.3	64.1	24.4	63.4	60.4	62.4	58.7	37.7	68.5	31.6	71.5	44.0	56.86
test	*Initial*	87.9	69.2	29.2	74.9	41.7	53.4	70.6	69.6	59.9	18.3	66.1	24.9	62.5	63.3	68.8	55.4	33.7	63.8	18.6	64.3	44.9	54.35
	EM-*Iter1*	88.2	70.3	30.0	75.3	45.2	58.3	72.0	72.0	66.8	19.0	67.4	29.8	67.6	65.1	72.2	57.9	43.2	69.2	23.9	65.4	44.5	57.31
	EM-*Iter2*	88.2	70.8	29.5	72.3	45.3	57.3	72.8	72.1	69.0	19.8	64.4	34.6	67.4	63.5	72.4	59.1	45.8	69.3	28.1	64.3	45.2	57.67

Table 5: Experimental results on PASCAL VOC validation and test set using the *Max Heuristic*, with a threshold η of (0.1, 0) for EM *Iteration 1* and EM *Iteration 2*, respectively.

Data	Stage	bkg	plane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mIoU
val	*Initial*	87.1	74.7	29.0	69.8	55.8	55.6	73.3	65.2	63.4	15.8	61.5	15.9	60.0	56.4	57.5	53.7	32.9	65.6	23.9	64.6	42.2	53.53
	EM-*Iter1*	87.6	74.3	29.7	69.9	58.3	57.1	76.3	68.0	68.7	17.7	62.9	19.8	64.1	59.6	61.2	57.2	37.3	67.3	27.6	69.8	43.5	56.09
	EM-*Iter2*	87.8	72.4	28.7	67.9	58.8	55.8	78.0	69.7	70.2	17.8	63.3	23.2	65.7	60.5	63.1	58.7	40.0	68.2	28.9	70.9	45.5	56.91
test	*Initial*	87.9	69.2	29.2	74.9	41.7	53.4	70.6	69.6	59.9	18.3	66.1	24.9	62.5	63.3	68.8	55.4	33.7	63.8	18.6	64.3	44.9	54.35
	EM-*Iter1*	88.1	70.4	30.2	75.4	46.4	56.3	73.3	71.7	67.3	19.6	66.4	29.9	67.3	63.4	72.4	58.2	42.9	70.8	22.6	64.7	46.1	57.31
	EM-*Iter2*	88.2	69.5	29.7	72.2	45.1	57.3	73.2	72.7	69.3	20.5	65.4	33.5	67.8	64.0	72.3	58.9	45.5	69.8	26.8	63.8	46.8	57.74

Table 6: Experimental results on PASCAL VOC validation and test using *Relative Heuristic*, with a threshold η of (0.05, 0.05) for EM *Iteration 1* and EM *Iteration 2*, respectively.

7.2. Reasons behind not using CRFs

Table 7 provides a more detailed version of the Table 2 shown in the main paper, highlighting the datasets and levels of supervisions required by the dependencies of the state-of-the-art weakly/semi supervised methods. Of the dependencies listed, we bring particular focus to the way CRFs have been used in the case of weakly supervised semantic segmentation problems where the CRFs hyper-parameters used are normally cross validated over a fully supervised pixel wise segmentation dataset. We, however, do not employ a CRF at any point in our pipeline for several reasons — (1) CRF hyper-parameters are generally fine-tuned on pixel-level annotated segmentations which makes the method not fully weakly-supervised. (2) Our EM algorithm relies on predicting a latent segmentation which is then treated as the ground-truth in order to update the parameters. Employing a CRF at this stage may enforce smoothness in the segmentation which may lead to noisy predictions. (3) The hyper-parameters of the CRF are incredibly sensitive, and if, for example, we wish to extend our framework to incorporate the learning of new object categories, this would require a fully-supervised dataset of the new categories along with the old ones for the cross-validation of the CRF hyper-parameters. This makes the approach highly non-scalable, and does not agree with the overall motivation for doing the weakly-supervised learning. Along with CRFs, notice that MCG is also trained over fully supervised semantic segmentation dataset.

The two dependencies used by our method are the saliency network of [23] and the attention network of [39]. In training the saliency network, supervision is provided by saliency masks. These are class-agnostic foreground masks of the salient instance/s in an image. We consider this a suitably weak form of supervision since once trained, the network can generalise to any semantic object category. In training the attention network, supervision is provided by image labels, which aligns with our weakly-supervised setting.

7.3. More examples of success and failure cases

Figure 4 and Figure 5 show several more qualitative examples of failure and success cases of our weakly-supervised semantic segmentation method. In terms of failure cases, our method struggles to segment thin or fine structures, for example, the wheels of bicycles and the legs of chairs and tables. Conventionally, the application of a CRF either in the training or as a post-processing step helps to better delineate such structures. However, for the above-mentioned reasons, we do not employ a CRF in our pipeline. Another common failure arises in the case of overlapping objects and less salient objects such as 'plant' and 'sofa'. Again, this can largely be attributed to the failure of the saliency and the attention networks, where the overlapped object is likely overlooked or considered less salient or attentive when compared with other objects present in the image. In terms of success cases, Figure 5 shows the progression of the segmentation from the *Initial Model* prediction to the final iteration of the EM algorithm. In these cases we can clearly see the improvement of the segmentations with each iteration, suggesting that the iterative nature of the EM algorithm and our learning objective are well suited to the weakly supervised segmentation task.

Dependency	Dataset	Supervision
Class Size	$D(P)$	Image labels + Bboxes
Saliency	$D(I+P)$	[34] Image labels
	MSRA-B [24]	[17] Bboxes
	MSRA10K [8] and DUT-OMRON [37]	[23] Saliency masks
Attention [39]	$D(I+P07+C)$	Image labels
Objectness [2]	Not $D(P)$	Image labels + Bboxes
Localization [41]	$D(I+P)$	Image labels
Superpixels [13]	NA	None
BBox BING [7]	$D(P07)$	Bboxes
MCG [3]	$D(P)$, BSDS [26]	Pixel labels
SS [35]	$D(P07)$	Bboxes
CRF [19]	–	Pixel labels (parameter cross-val)

Table 7: **Dependency Table**. This table represents the degree of supervisions required by different dependencies of existing weakly-supervised methods for the *semantic segmentation task*. The dataset notation is similar to the one explained in the main paper. The extra notations, $P07$ and C represents the PASCAL VOC 2007 and Microsoft COCO [22] datasets, respectively.

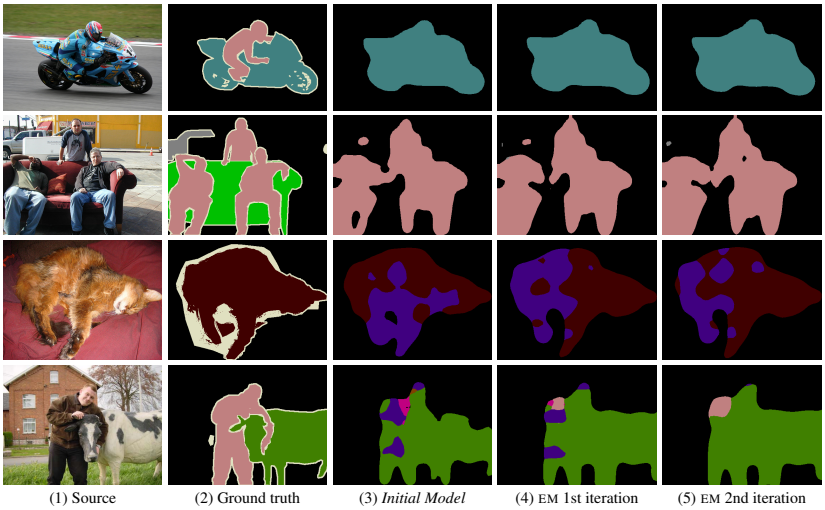

Figure 4: **Failure cases.** Visual results for the weakly-supervised semantic segmentation using our proposed EM based method, from the *Initial Model* (3rd column) to the final iteration of the EM algorithm (5th column).

References

[1] F. Ahmed, D. Tarlow, and D. Batra. Optimizing expected intersection-over-union with candidate-constrained crfs. In *ICCV*, 2015. 2, 5

[2] B. Alexe, T. Deselares, and V. Ferrari. Measuring the objectness of image windows. In *PAMI*, 2012. 6, 7, 12

[3] P. Arbelaez, J. Pont-Tuset, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping. In *CVPR*, 2014. 2, 6, 7, 12

[4] A. Bearman, O. Russakovsky, V. Ferrari, and F.-F. Li. What's the point: Semantic segmentation with point supervision. In *ECCV*, 2016. 1, 2, 3, 5, 7

[5] S. Chandra and I. Kokkinos. Fast, exact and multi-scale inference for semantic image segmentation with deep gaussian crfs. In *ECCV*, 2016. 1

[6] L.-G. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected. In *ICLR*, 2015. 1, 6

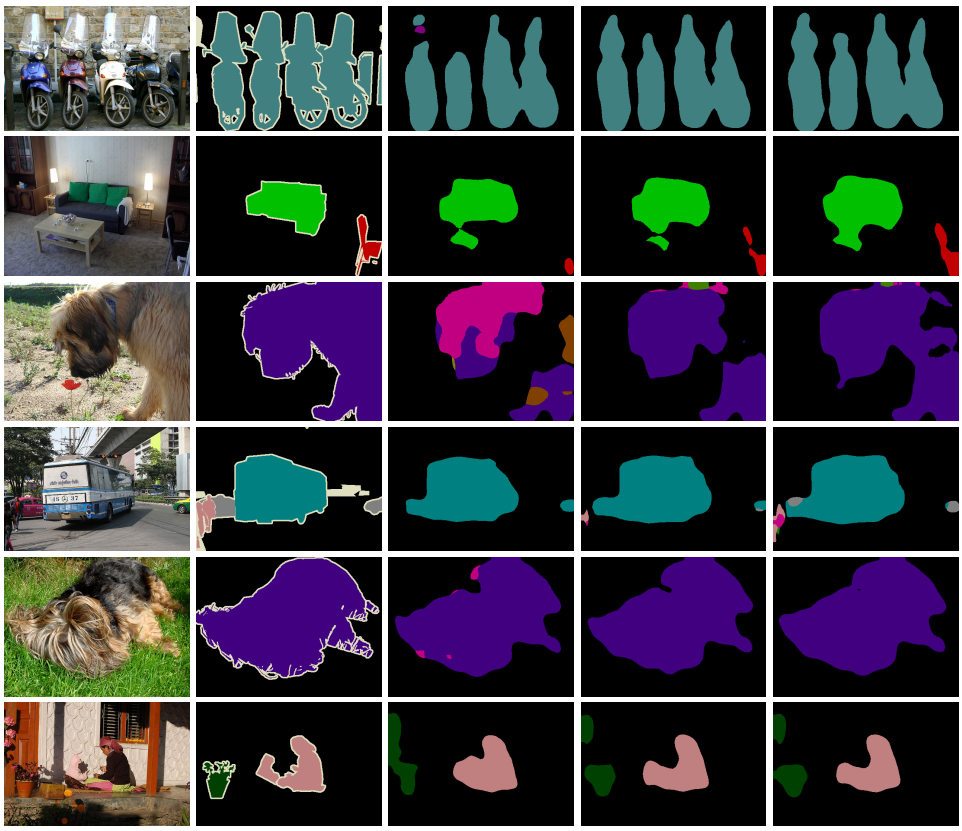

Figure 5: **Success cases.** Visual results for the weakly-supervised semantic segmentation using our proposed EM based method. It is clear that as we move from the *Initial Model* (3rd column) to the second iteration of the EM algorithm (5th column), the segmentation quality improves. This shows that each EM iteration is learning semantically meaningful information using only image-level labels.

[7] M. Cheng, Z. Zhang, W. Lin, and P. H. S. Torr. Bing: Binarized normed gradients for objectness estimation at 300fps. In *CVPR*, 2014. 6, 7, 12

[8] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. S. Torr, and S.-M. Hu. Global contrast based salient region detection. *IEEE TPAMI*, 37(3):569–582, 2015. 5, 12

[9] M. Cogswell, X. Lin, S. Purushwalkam, and D. Batra. Combining the best of graphical models and convnets for semantic segmentation. In *arXiv:1412.4313*, 2014. 2, 5

[10] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. In *Journal of the Royal Statistical Society*, 1977. 1, 3

[11] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 2, 5

[12] M. Everingham, S. M. A. Eslami, L. V. Gool, C. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge a retrospective. In *IJCV*, 2014. 2, 5

[13] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph based image segmentation. In *IJCV*, 2004. 6, 7, 12

[14] B. Hariharan, P. Arbelaez, L. Bourdev, S. Maji, and J. Malik. Semantic contours from inverse detectors. In *ICCV*, 2011. 6

[15] C. F. Jeff Wu. On the convergence properties of the em algorithm. In *The Annals of Statistics*, 1983. 1, 5

[16] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM International Conference on Multimedia*, 2014. 6

[17] H. Jiang, J. Wang, Z. Yuan, Y. Wu, Z. N., and S. Li. Salient object detection: A discriminative regional feature integration approach. In *CVPR*, 2013. 6, 7, 12

[18] A. Kolesnikov and C. H. Lampert. Seed, Expand and Constrain: Three Principles for Weakly-Supervised Image Segmentation. In *ECCV*, 2016. 1, 2, 6, 7, 10

[19] P. Krahenbuhl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011. 6, 7, 12

[20] S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 1951. 2

[21] D. Lin, J. Dai, J. Jia, K. He, and J. Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, 2016. 1, 2, 3

[22] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 1, 12

[23] N. Liu and J. Han. Dhsnet: Deep hierarchical saliency network for salient object detection. In *CVPR*, 2016. 2, 5, 6, 7, 11, 12

[24] T. Liu, J. Sun, N. Zheng, X. Tang, and H. Shum. Learning to detect a salient object. In *CVPR*, 2007. 12

[25] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1, 6

[26] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, 2001. 12

[27] G. J. McLachlan and T. Krishnan. *The EM algorithm and extensions*. Wiley, 1997. 1, 3

[28] S. Nowozin. Optimal decisions from probabilistic models: the intersection-over-union case. In *CVPR*, 2014. 2, 5

[29] G. Papandreou, L.-C. Chen, K. P. Murphy, and A. L. Yuille. Weakly- and semi-supervised learning of a DCNN for semantic image segmentation. In *ICCV*, 2015. 1, 2, 3, 4, 5, 6, 7

[30] D. Pathak, P. Krahenbuhl, and T. Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *ICCV*, 2015. 1, 2, 7, 10

[31] D. Pathak, E. Shelhamer, J. Long, and T. Darrell. Fully convolutional multi-class multiple instance learning. In *ICLR*, 2014. 2, 10

[32] P. O. Pinheiro and R. Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, 2015. 1, 2, 7

[33] X. Qi, Z. Liu, J. Shi, H. Zhao, and J. Jia. Augmented feedback in semantic segmentation under image level supervision. In *ECCV*, 2016. 1, 2, 7, 10

[34] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR*, 2014. 6, 7, 12

[35] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders. Selective search for object recognition. In *IJCV*, 2013. 6, 7, 12

[36] J. Xu, A. Schwing, and R. Urtasun. Learning to segment under various forms of weak supervision. In *CVPR*, 2015. 1, 2, 3

[37] C. Yang, L. Zhang, H. Lu, X. Ruan, and M. H. Yang. Saliency detection via graph-based manifold ranking. In *CVPR*, 2013. 12

[38] W. Yunchao, L. Xiaodan, C. Yunpeng, S. Xiaohui, M.-M. Cheng, Z. Yao, and Y. Shuicheng. STC: A simple to complex framework for weakly-supervised semantic segmentation. In *arXiv:1509.03150*, 2015. 1, 2, 3, 6, 7

[39] J. Zhang, Z. Lin, J. Brandt, X. Shen, and S. Sclaroff. Top-down neural attention by excitation backprop. In *ECCV*, 2016. 2, 3, 5, 6, 7, 11, 12

[40] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. Torr. Conditional random fields as recurrent neural networks. In *ICCV*, 2015. 1

[41] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 6, 7, 12